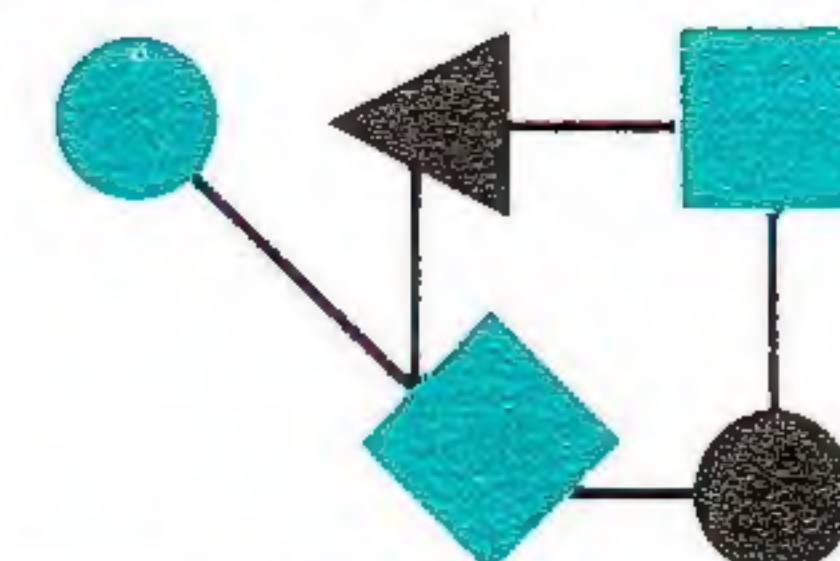# CONNEXIONS

The Interoperability Report

May 1993   *Special Issue: Focus on Europe*   Volume 7,   No. 5

*ConneXions —
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

From the Editor

This month we take a look at the state of networking in Europe. A great deal of Internet growth has taken place in this region over the last few years, and several major developments are currently underway. As you probably already know, the first European INTEROP will take place in Paris, October 25–29, 1993. This issue offers a glimpse of some of the topics that will be discussed at the Paris conference, but I must emphasize that the articles herein by no means give a complete picture. In fact, I'd like to use this opportunity to solicit more articles from Europe for the October issue. Please contact connexions@interop.com for author guidelines and information about deadlines. I look forward to hearing from you.

Our first article describes EBONE, the European Internet backbone. This backbone was put together in record time, and is an initiative to promote close collaboration among existing European network service providers to provide a general purpose pan-European IP backbone. The article is by Bernhard Stockman.

The business and research case for pan-European Broadband networks is described next in an article by Brian Carpenter. The article outlines several technological and political factors that both enable and hinder the deployment of such networks in Europe. It should be noted that local telecommunications regulations, tariffs, protocol mandates and "the installed base" vary greatly from country to country, and this makes the realization of multi-national networks a significant challenge.

Most traditional data networks in Europe were based on X.25, provided typically by the national PTT, but in some cases operated privately using leased lines. In the UK, the *Joint Academic Network* (JANET) has been in operation since the 1970s. JANET has undergone many changes and upgrades over the years as outlined in an article by Jon Crowcroft on page 14.

Another European tradition is the promotion of OSI. While many of today's European networks use TCP/IP, it is still the case that more attention has been paid to OSI—both in terms of research and available products—as compared to the United States. Our final article, also from the UK, is a true "report from the trenches." Paul Barker and Colin Robbins report on their efforts to demonstrate OSI software at a number of exhibitions, most of them using X.25 as the underlying network service. In their experience, many X.25 implementations are of poor quality when compared with other networking technology, and they suggest that the cause of OSI may be better served by using non-OSI networking technology instead of X.25. Paul and Colin have promised a follow-up article which we hope to publish in the October issue.

EBONE, The European Internet Backbone

by Bernhard Stockman, SUNET

Abstract

The requirement for Europe-wide network connectivity at increasing bandwidth and for high-speed connectivity to non-European networks has long been recognized. In recent years, this demand has focused on IP services. Until late 1991 such connectivity was available only to individual national and international research networks. EBONE is an initiative to promote close collaboration among existing European network service providers to provide a general purpose pan-European backbone. This article gives a report of the status and the near future development of EBONE.

Introduction

EBONE began in September 1991 when representatives of several European academic and research networks met to resolve long-standing European connectivity problems. Their approach was to evaluate existing available links, to look for opportunities to bring these links together quickly under a unified approach, and to make plans to enhance these links. This was documented in the initial EBONE proposal. [1]

Contributions were secured, a management structure was established, operational procedures were put in place, and an overall contribution-oriented funding approach was agreed. Each participating organization has signed a *Memorandum of Understanding* which defines the terms of EBONE membership and the resources which each member contributes to the EBONE effort. [2]

EBONE focuses on supporting networking organizations which serve the European academic and research communities. Through EBONE, European researchers have improved access and higher-performance connections to their colleagues throughout Europe and to other continents.

Furthermore, by encouraging the participation of commercial network service providers (e.g., PTTs, information technology companies), EBONE will increase the size of the participating communities, reduce individual costs, encourage the participation of industrial researchers, and stimulate the creation of competitive international IP networking services in Europe.

EBONE organization

The EBONE organization is formed by its members, currently 21 national and international networks from the R&D and commercial sectors. The EBONE members nominates its *Management Committee* and nominates participation in the technical subcommittee, the *EBONE Action Team*.

EBONE is operated by the EBONE Operations Team consisting of operational managers from each of the EBS sites plus the *EBONE Network Operations Center* at The Royal Institute of Technology (KTH) in Stockholm.

During 1992 EBONE was financed by contributions from the members. This was a pragmatic approach to get EBONE up and running with minimal delays. After this, the principle of fair sharing of costs has been adopted. The total cost for EBONE is distributed to its members based on the connected bandwidth. The 1993 budget is close to 3 million ECU (around the same amount in US dollars). Of these costs around 2 million ECU (65%) is used for the leased lines and 300 kECU (11%) for the EBONE staff.

The total cost is distributed to the EBONE members based on connected capacity as counted in 64Kbps equivalents. For 1993 the cost for a 64Kbps equivalent is 64 kECU. A weighted system, based on the PTT tariffs for fractional E1 services, is used to calculate costs for connections at higher capacities giving proportionally a lower connection cost at higher capacities.

Functional description

EBONE today operates a core backbone between Amsterdam, Bonn, Geneva, London, Paris and Stockholm. Regional, national and international networks connect to these core backbone sites. The links interconnecting the EBONE core sites operates at speeds between 256 and 512Kbps. Recently it was decided to upgrade the Stockholm–Amsterdam–Geneva–Paris path to T1/E1. Intercontinental links are currently connected at Bonn (512Kbps), Geneva (1544Kbps), London (1024Kbps), Paris (1544Kbps) and Stockholm (1544Kbps).

A three level hierarchy is imposed with the EBONE kernel at the top interconnecting regional networks which in turn connect sites within the regional networks. The EBONE kernel does not enforce any restrictions on the traffic being forwarded to and from EBONE contributing organization, as long as this traffic is not regarded as harmful to the overall EBONE functionality. Internet IP is the supported layer three production protocol. ISO CLNS is provided as a pilot for the RARE/COSINE CLNS Project.

EBONE kernel

The EBONE core sites known as *EBONE Boundary Systems,* is where the traffic forwarding decisions are made and the backbone lines interconnecting the EBSs are terminated. The *EBS Technical Specification* document gives recommendations for EBS installations. [3]

For reasons of resilience, each EBS has at least two EBONE core links, connecting into two other EBSs and these links shall, as far as possible, be routed in separated paths between the EBSs.
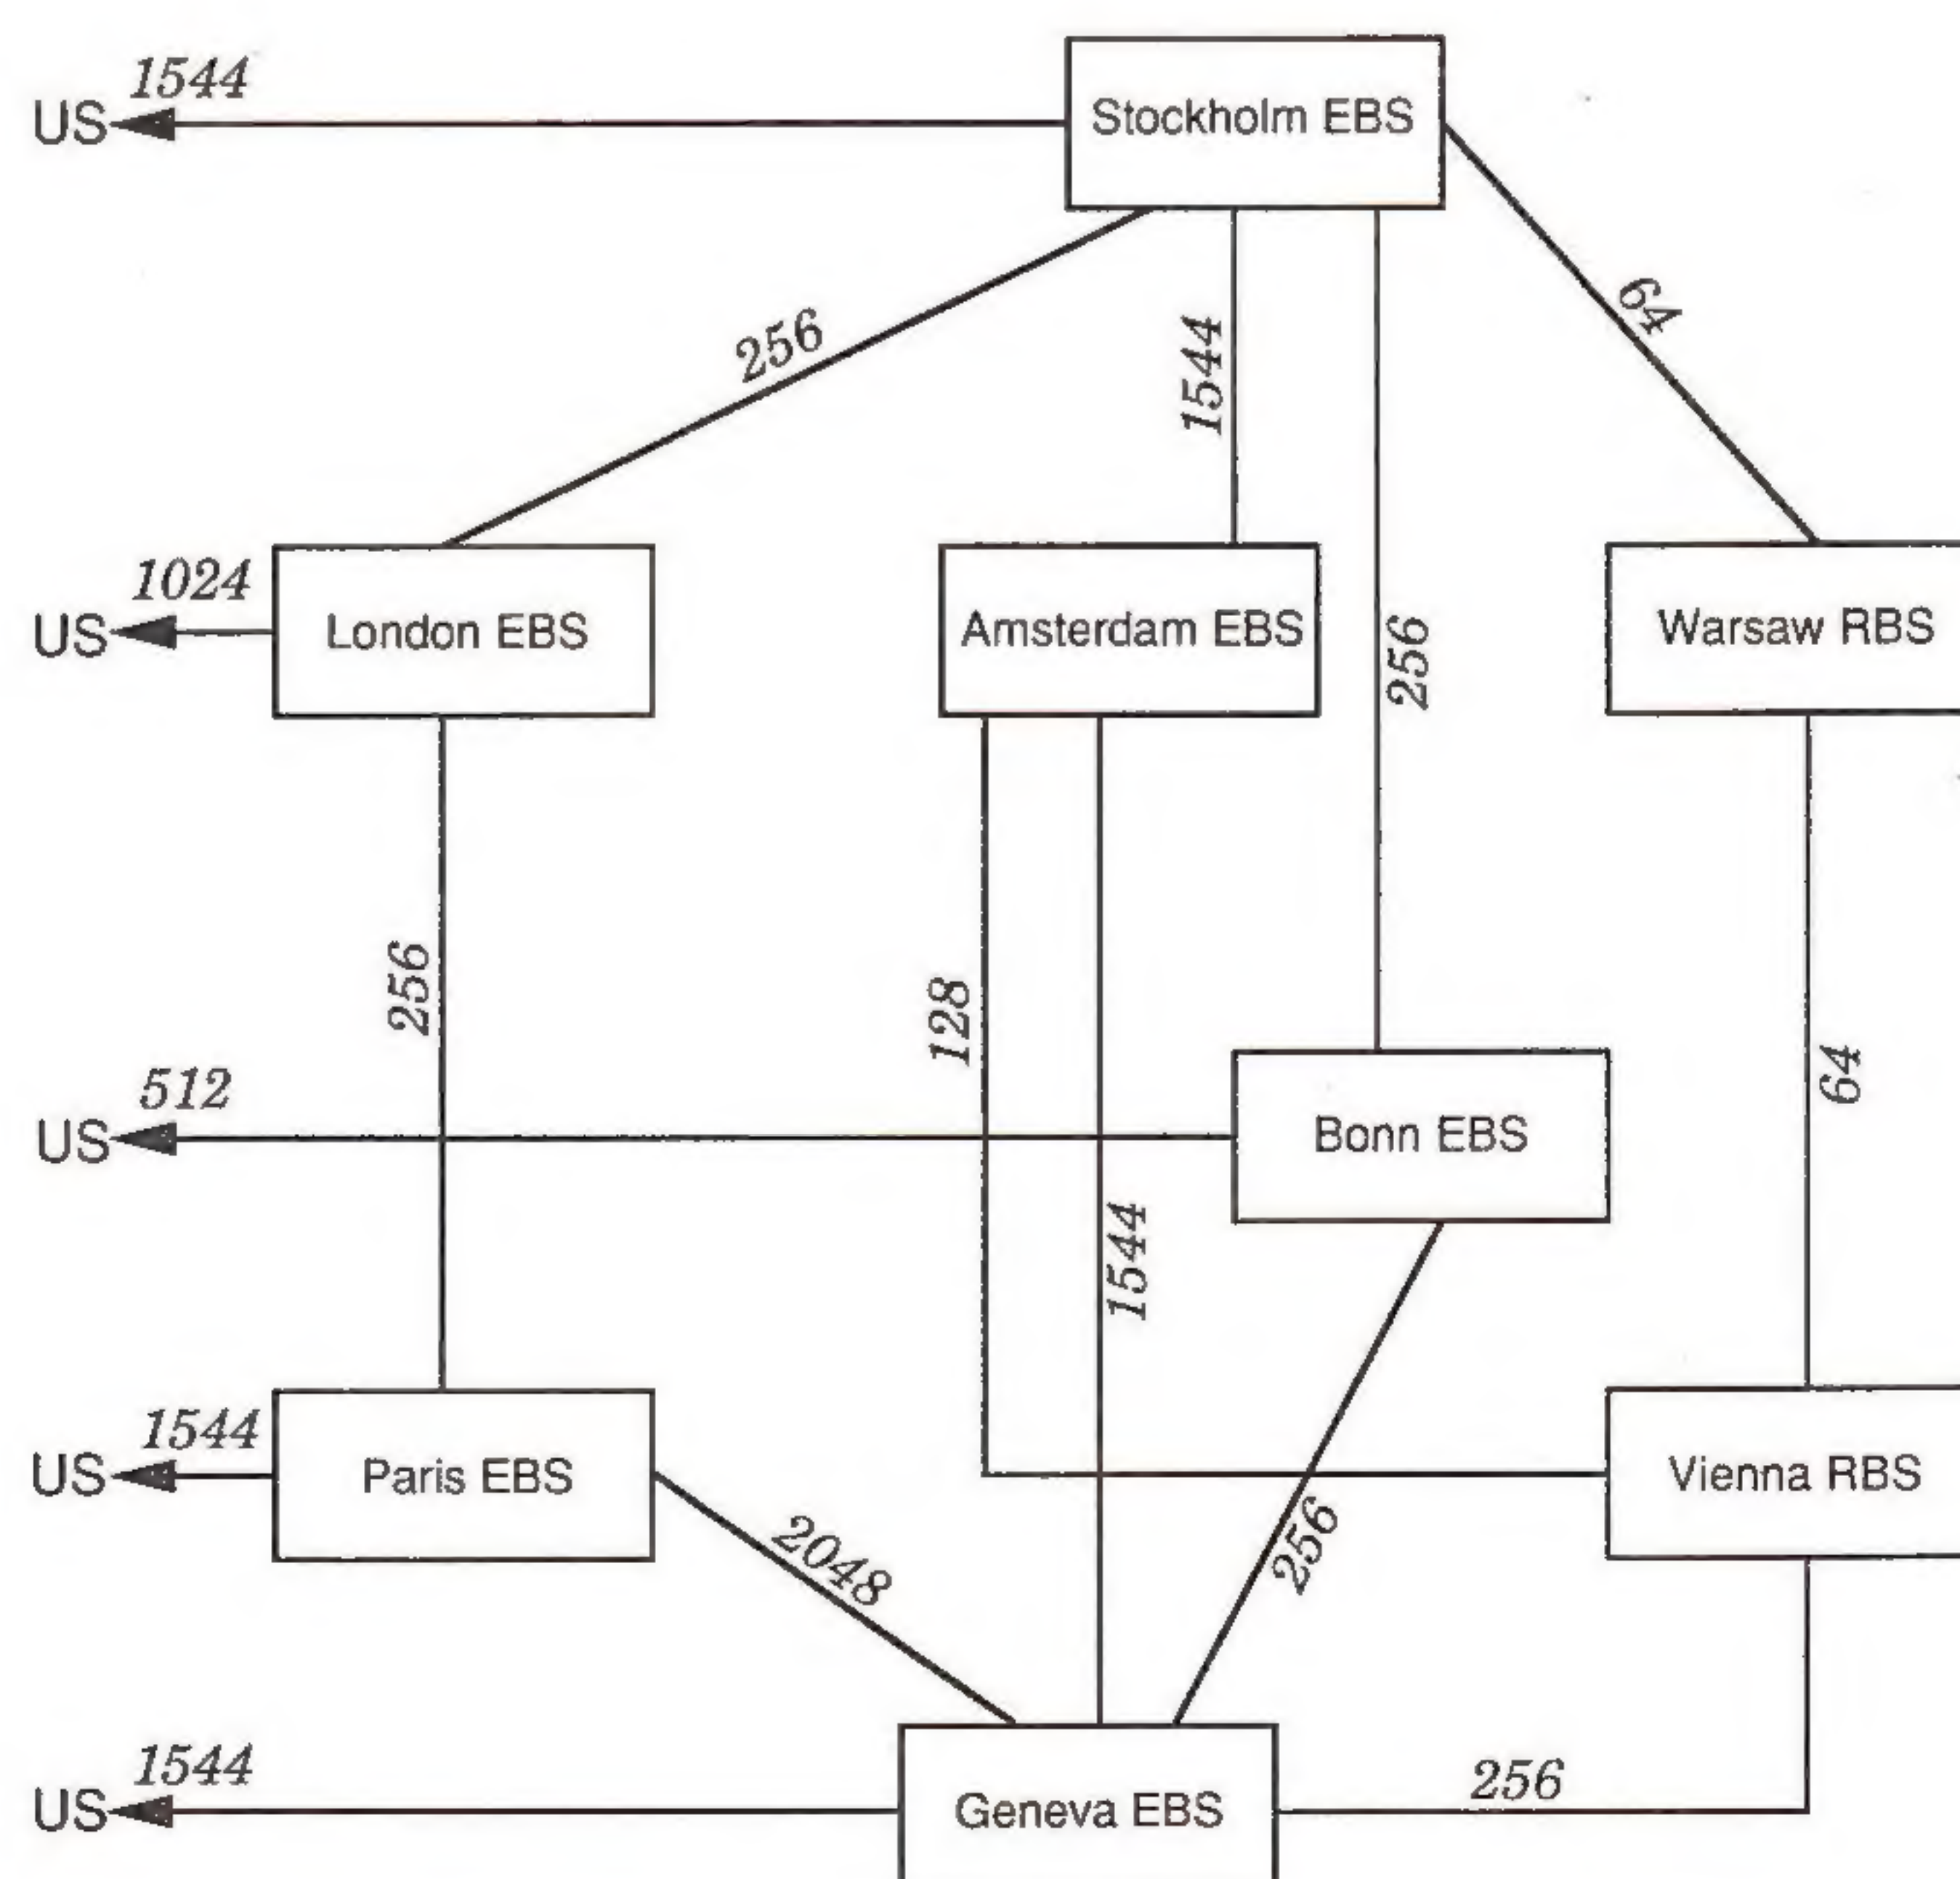

Figure 1: EBONE Basic Topology, March 1993

Note: In the above map the dual homed regions connected in Vienna (ACOnet) and in Warsaw (NASK) are included although not formally part of the EBONE core. Discussions are, however, ongoing for the possible integration of these connections into the EBONE core.

3

EBONE *(continued)*

Connections between EBONE and Regional Networks

The boundary between the EBONE kernel and connecting networks is described in terms of the functionality of the systems at either side and the specifications of their common interface. The defined functions are handled by the *EBONE Boundary System* (EBS) and the *Regional Boundary System* (RBS).

Regional networks connect to the EBONE via Regional Boundary Systems (RBS). Regional policies are implemented in the RBSs. This makes it possible for each connecting network to enforce their local policy on the traffic they accept without restricting the traffic possible to forward in the EBONE kernel.
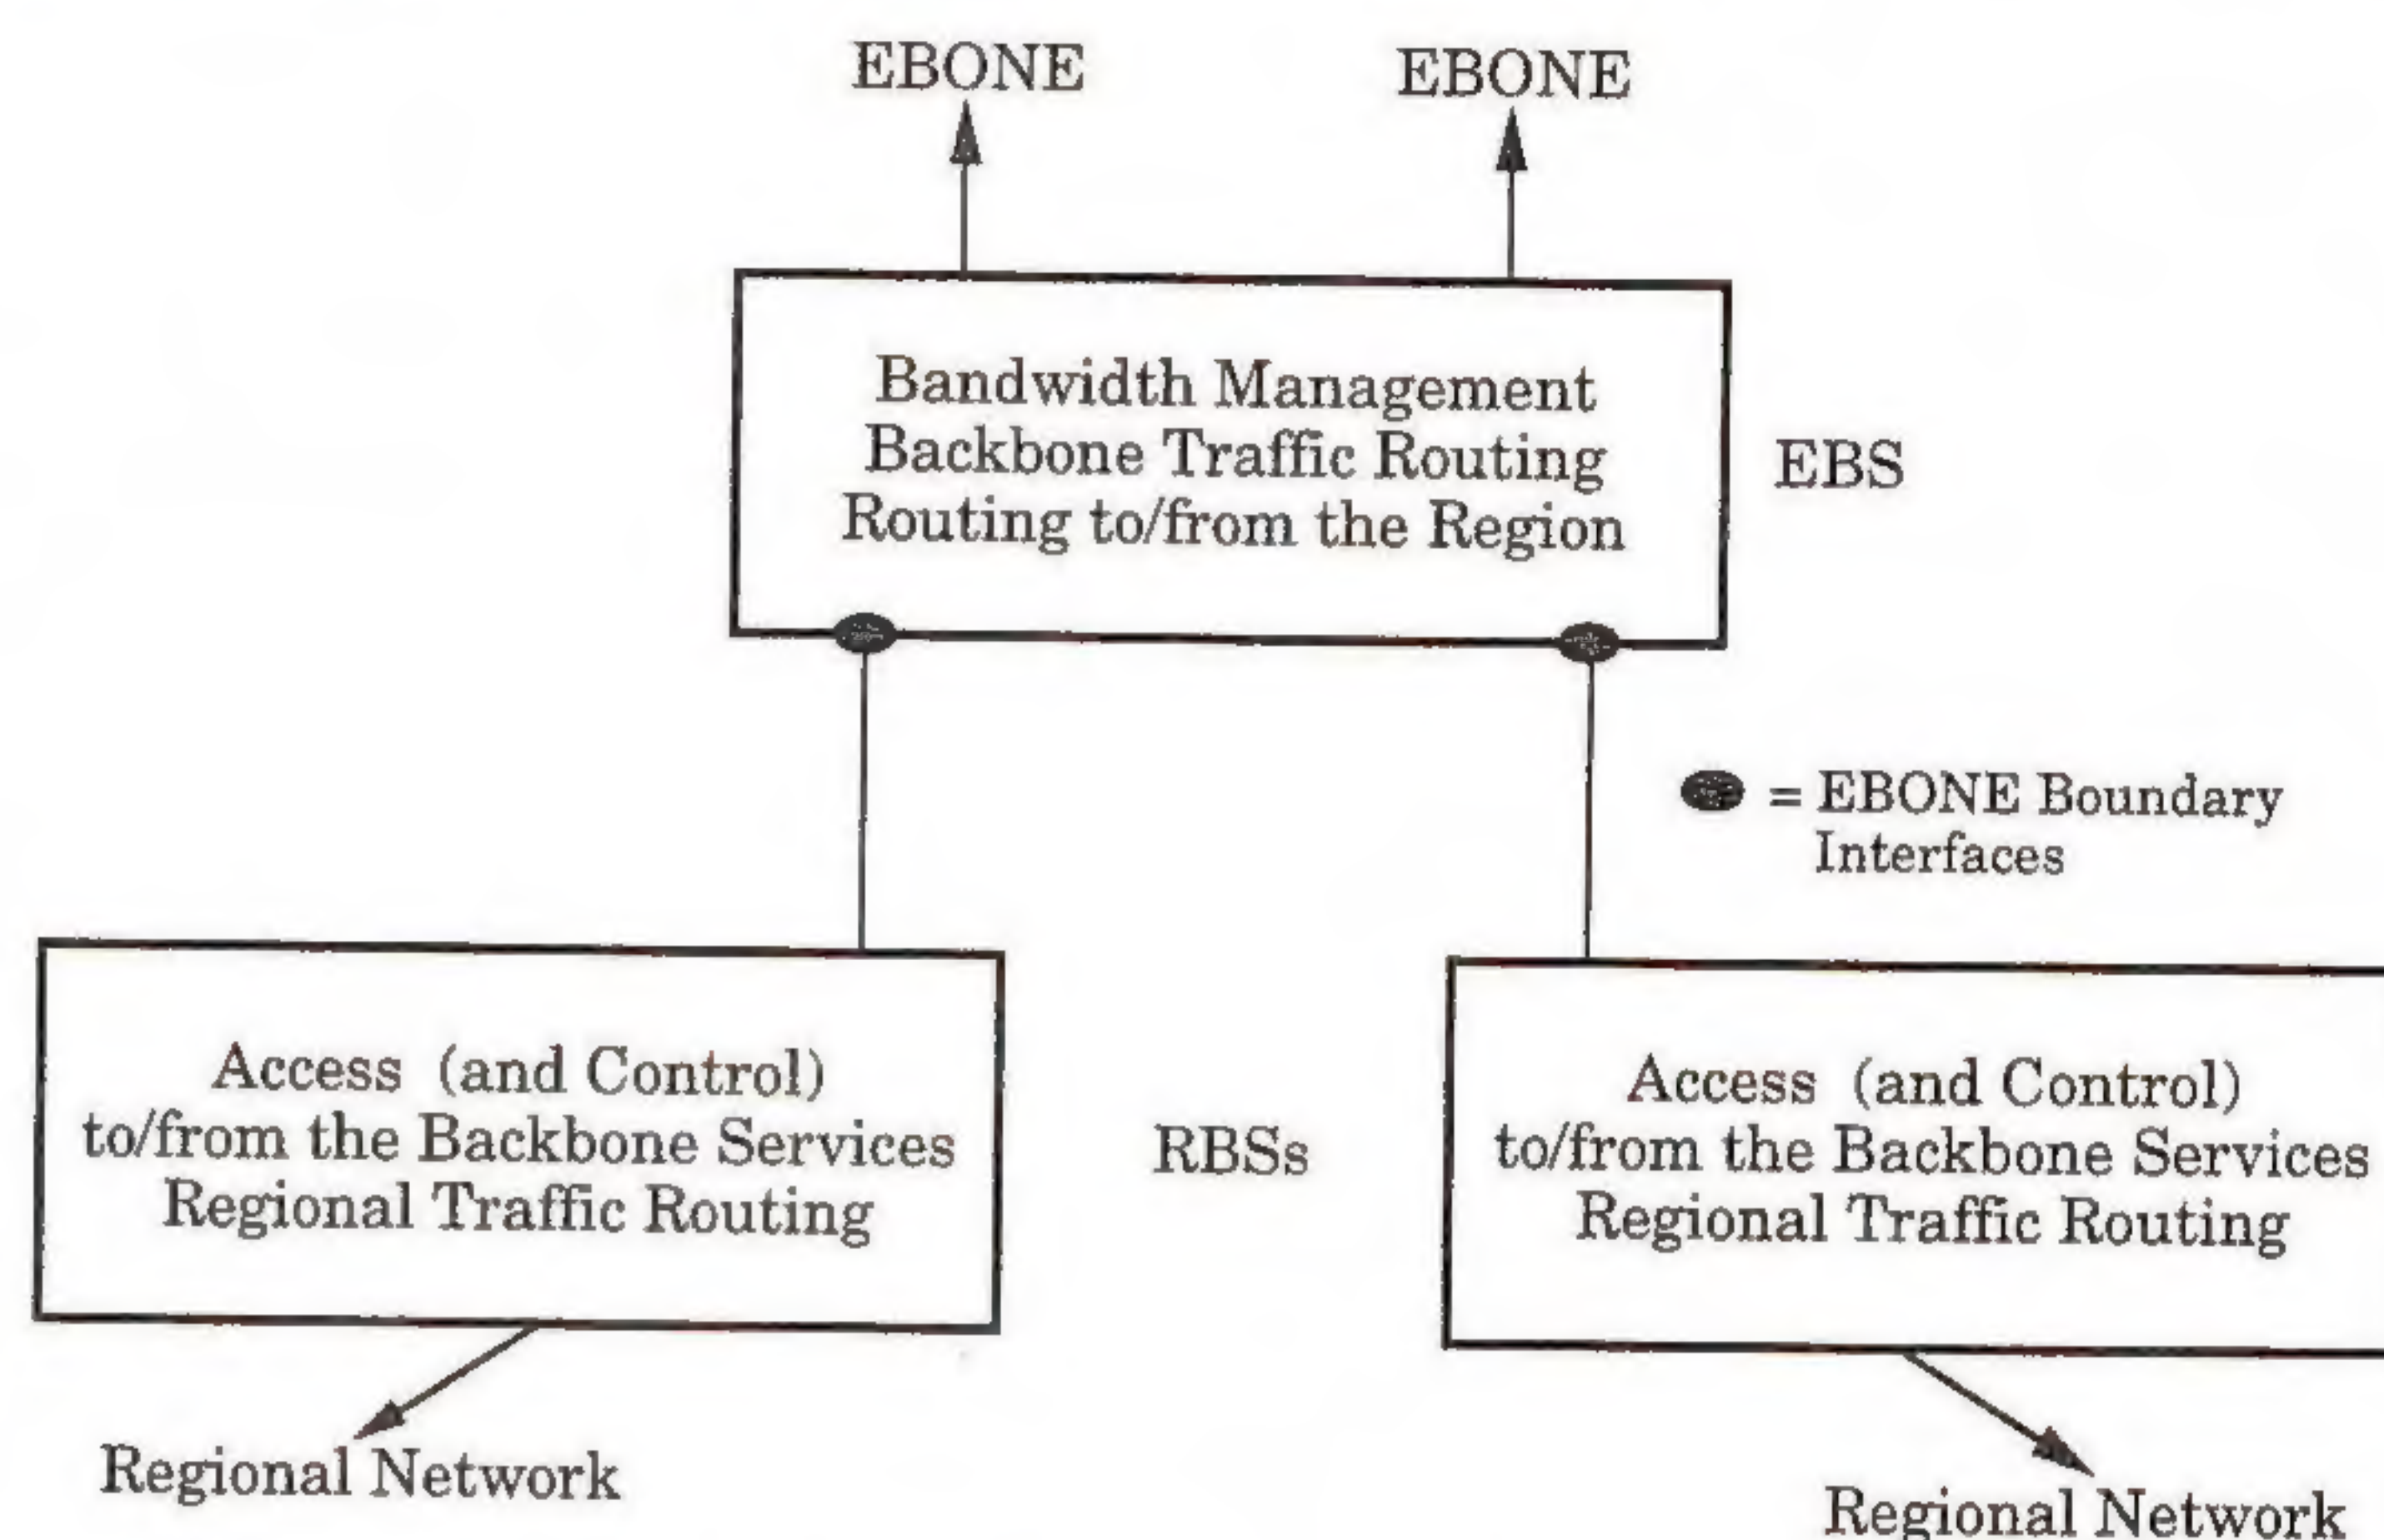

Figure 2: Regional connectivity to the EBONE

Both EBS and RBS shall be implemented using similar technology to assure interoperability with regards to forwarding of traffic and management of the EBSs and RBSs. For further details on the implementation of the initial backbone see [4].

As of March 1, 1993 the following regional networks are connected via an EBS–RBS connection.

Amsterdam EBS:
SURFnet (Netherlands)
ACOnet (Austria)
RedIRIS (Spain)
ECRC (Germany)
EUnet (Europe)

London EBS:
JANET (United Kingdom)
PIPEX (United Kingdom)
HEANET (Ireland)

Bonn EBS:
GMD (Germany)
Geneva EBS:
CERN (Europe)
ACOnet (Austria)
ARIADNet (Greece)
ILAN (Israel)
EARN (Europe)
SWITCH (Switzerland)

Paris EBS:
RENATER (France)
BELNET (Belgium)
FORTH (Greece)
RCCN (Portugal)

Stockholm EBS:
NORDUnet (North Europe)
SWIPNet (Sweden)
TELECOM Finland
TIPNET (Sweden)

Negotiations are ongoing with the CEC-funded *EuropaNet* for the provision of intercontinental access to EuropaNet via EBONE. Austria recently signed a Memorandum of Understanding with several central and east European countries where it was agreed that an Austrian EBS would be the connection to EBONE for these countries and that ACOnet was given a mandate to negotiate with EBONE on behalf of these countries.

Several efforts are ongoing to connect the CIS states to EBONE. The Baltic countries are today connected to NORDUnet and by that to EBONE. There is a big interest in establishing connectivity from Europe and the USA to Russia, and discussions are now happening for how this shall be realized. As CoCom regulations are still in effect, there are problems in deploying adequate technology in the CIS countries for the establishment of efficient network connectivity to the rest of the global Internet.

Routing implementation

Routing within the EBONE kernel is effectively "open" from the point of view that no access control will be put on EBS routing announcements either in-bound or out-bound to the EBONE kernel. For control purposes and pragmatic reasons, access control will take place at the EBS–RBS interface based on information in the RIPE database. The RIPE database routing object is currently being enhanced to include AS based policy information [10].

It is possible to use BGP for Intra-Autonomous System (AS) routing to maintain a coherent view of BGP derived routes. This is known as *Internal BGP* (IBGP). As EBONE uses BGP3 as EGP, IBGP can be used together with IGP within the backbone. This will give added functionality of avoiding routing loops and distribute AS path information within the backbone. Practically this is done by having a full mesh of IBGP sessions between all EBS systems. Next hop information within the backbone is today derived using the IGRP, but discussions are ongoing for the changing to OSPF or IS–IS.

Management and operations

The management and operations of the EBONE is undertaken by the *EBONE Operations Team* (EOT). The EOT is formed of the set of EBS managers and the EBONE NOC. Each EBS site must have one designated EBS manager. In order to coordinate EBS management, an EBONE *Network Operations Center* (NOC) has been established at the Royal Institute of Technology (KTH) in Stockholm, Sweden.

The EBS manager is responsible for the local operation of his EBS installation and acts as a contact point between the EBONE NOC and the EBS installation, and as a contact point between the EBS installation and regional connections to that EBS. Fault reporting from a connected network shall in the first instance happen to the relevant EBS site which, in collaboration with the NOC, will be responsible for fault identification and repair.

As EBONE management is distributed it is important to have a security scheme in place for accessing the EBS routers. This has been implemented using TACACS servers. When operating and managing several routers in such a collaborative way, the consistency of the access database is very important. Eventually an automated scheme for the update mechanism should be available to take care of this.

The EBONE NOC is, besides doing daily operational coordination, responsible for producing monthly reports on traffic statistics and other significant events in the backbone. Each EBS site collects statistics for its local EBS.

EBONE *(continued)*

The processed data is then distributed from the EBS sites to the NOC for inclusion in the EBONE reports. It is recommended that the data is stored according to the current IETF OPSTAT recommendation [11]. As soon as there exists a publicly available statistics collection package using the OPSTAT format, each EBS should run this.

For a well run and well maintained network, a "Trouble Ticket" system is essential. Currently, there appears to be a lack of a good generic trouble ticket system available. However, several EBS sites have already adopted their own somewhat ad-hoc systems. This is not a problem providing they conform to the same standard presentation format as described in the RIPE recommendation on operational contacts [12]. A trouble-ticket research group has been set up within the EOT to review this matter in more depth. For a full description of the management and operations see [6] and [7].

Utilization of the EBONE core

From the beginning, the EBONE core lines have been more or less fully utilized. Utilization figures around 80 percent of full capacity is the rule during office-hours. The path Stockholm–Amsterdam–Geneva–Paris has had severe performance degradation due to overload. For this reason, the EBONE Management decided recently, after an additional funding was made available by RENATER, SURFnet and NORDUnet, to upgrade this path to 1544/2048Kbps as indicated in Figure 1.

One principle in the usage of the EBONE inter-continental lines is that these lines should be possible to use for mutual backups. This scheme requires, however, that the EBONE core lines must be on par with the inter-continental capacity so as not to form bottlenecks for backup traffic. The EBONE core lines to the United Kingdom (256 Kbps) are too low in capacity if being used in this way and resources have to be found for upgrading.

Statistical measurements of the various EBSs are not yet fully coordinated but efforts in doing so are ongoing. The staff at the Amsterdam EBS will do the necessary coordination. As IP accounting on the routers is necessary for getting the full picture of the network–network and RBS–RBS traffic, work has been started in gathering such statistics. This is, however, very much in the piloting phase and methods for evaluating and presenting such figures has to be worked out in more detail. Some statistics gathered from the Stockholm EBS are shown below:

Megabytes during January 1993:

	Input	Output
Stockholm–Amsterdam:	40635.2	53839.6
Stockholm–London:	17876.0	36020.7
Stockholm–US:	121632.6	123293.2

Utilization during January 1993:

	Input % aver	peak	Output % aver	peak
Stockholm–Amsterdam:	25.88	54.83	34.27	72.86
Stockholm–London:	22.91	67.49	46.18	85.04
Stockholm–US:	44.18	75.85	44.99	72.83

Kbps during January 1993:

	Input		Output		Max
	aver	peak	aver	peak	speed
Stockholm–Amsterdam:	133	283	176	387	512
Stockholm–London:	59	173	118	220	256
Stockholm–US:	401	721	407	664	768

Connectivity to the Global Internet

For reasons of full connectivity to the global Internet, and for reasons of routing efficiency, all EBONE inter-continental lines are connected to the *Global Internet eXchange* (GIX) in Washington, DC. A proto-GIX is implemented on the *Metropolitan Area Ethernet East* (MAE-East) which provides Ethernet capacity between some POPs in the DC area.

The initial work with the GIX started a couple of years ago and has progressed within the *Intercontinental Engineering and Planning Group* (IEPG), a technical subcommittee to the *Consultative Committee for Intercontinental Research Networking* (CCIRN). See [8] for a discussion of the basic concepts.

Since the first GIX discussion, the EBONE has developed and by that the need for coordinated routing between EBONE and the rest of the global Internet. A project has for this reason been initiated by SURFnet, funded via RARE and hosted at the RIPE NCC in Amsterdam. The project aims for a specification of a European route server providing a consistent view with respect to routing of traffic to European destinations. The ambition is to develop, pilot and deploy the European Route Server as part of the GIX.

The service consists of the route server [9] combined with the RIPE routing database [10]. Using the routing policy information stored in the RIPE database, it will be possible to compile this information into a configuration to be loaded in the route server. Networks peering with the European Route Server will thus be given an optimal path for forwarding traffic to European destinations.

EBONE evolution

Demand for EBONE services is growing rapidly. Additional links and increased capacity will be installed to satisfy this demand, and several such upgrades are under planning. This is especially the case for the fast growing need for connectivity to Central and East European Countries. With the advent of commercial IP providers on the European scene, the possibility of getting cheap capacity between EBSs increases. One way may be to let pan-European IP providers install capacity between EBSs and then use the EBSs as connectivity points bringing the global Internet to some central access points similar to what has been outlined for the NAPs within the USA.

With the coming of demanding applications like audio and video conferencing, remote visualization, interactive animation etc., there will be a need to ensure that the EBONE EBSs are interconnected with at least E3 (34Mbps) lines (the European PTT standard corresponding to the US T3 45Mbps). An interesting possibility currently being discussed within various groups is ATM. ATM pilots are today taking place on a national and regional scale in Europe and the need to go pan-European is not far away and has to be taken into consideration when planning for the near-future evolution of the EBONE.

EBONE *(continued)*

References

[1] "The EBONE-92 proposal," September 1991.

[2] "The EBONE Memorandum of Understanding," December 1991.

[3] Juha Heinanen, Peter Lothberg, Bernhard Stockman, "The EBONE Boundary System Technical Specification," February 8, 1992.

[4] Bernhard Stockman, "The EBONE Implementation Plan," February 24, 1992.

[5] Tony Bates, Peter Lothberg, "EBONE Routing and Addressing Plan," April 1992.

[6] Bernhard Stockman, "EBONE Management and Operations," April 1992.

[7] Tony Bates, "EBS Operational Recommendations," February 1993.

[8] Guy Almes, Peter Ford, and Peter Lothberg, "Proposal for Global Internet Connectivity," June 1992.

[9] Tony Bates, Daniel Karrenberg, Peter Lothberg, Bernhard Stockman and Marten Terpstra, "Internet Routing in a Multi Provider, Multi Path Open Environment," February 1993.

[10] Tony Bates, Jean-Michel Jouanigot, Daniel Karrenberg, Peter Lothberg and Marten Terpstra, "Representation of IP Routing Policies in the RIPE Database," March 1993.

[11] Bernhard Stockman, "A Model for Common Operational Statistics," RFC 1404, November 1992.

[12] Daniel Karrenberg, "Recommendation on Operational Contacts," April 1992.

[13] "Réseaux Associés pour la Recherche Européenne (RARE)," *ConneXions,* Volume 6, No. 1, January 1992.

[14] Goldstein, S. & Michau, C., "Convergence of European and North American Research and Academic Networking," *ConneXions,* Volume 5, No. 4, April 1991.

[15] Stockman, B., "Current Status on Networking in Europe," *ConneXions,* Volume 5, No. 7, July 1991.

[16] Crowcroft, J., "UK Internet Protocol Connectivity Activities," *ConneXions,* Volume 5, No. 6, June 1991.

[17] Crowcroft, J., "International Internetworking," *ConneXions,* Volume 2, No. 4, April 1988.

EBONE documents are available via anonymous FTP from host `archive.ripe.net` and `nic.nordu.net` in the directory `EBONE`. RIPE documents are available via anonymous FTP from host `archive.ripe.net` in the directory `/ripe/docs/ripe-docs`.

BERNHARD STOCKMAN graduated as Master of Science in Electric Engineering and Computer Systems from the Royal Institute of Technology in Stockholm Sweden 1986. After a couple of years as a researcher in distributed computer systems, he was employed by the NORDUNET and SUNET Network Operation Center. Today Bernhard is the European co-chair of the Intercontinental Engineering and Planning Group (IEPG) and involved in networking infrastructural questions both within Europe and on a global scale. He chairs the EBONE Action Team (EAT) responsible for the technical development of EBONE. His e-mail is: `boss@sunet.se`

The Importance of
Pan-European Broadband Networks

by Brian E. Carpenter, CERN

Introduction

A new and revolutionary generation of data networking technology and applications is imminent. In this article I will outline why this revolution is important (the business case), give a snapshot of the rapidly evolving technology, and finally comment on some essentially political questions concerning its deployment on a European scale.

The business case

I will return several times to the topic of money. This may seem strange, coming from a scientist, but the fact is that scientists find money to be just as fundamental an issue as do industrialists or businessmen. A research scientist's problem with money is a bit different from a business person's: the scientist fights for the biggest budget he or she can get, and then spends it in the most cost effective way. The business person can justify a bigger budget by the hope of bigger returns. A few years ago, CERN decided that we needed 2Mbps data connections to various laboratories in our member states, but the tariffs were well beyond the reach of our existing budget. I happened to ask a telecommunications manager from one of the main Swiss banks whether he also found the tariffs for international leased lines to be scandalously high; he simply said "Yes, but we pay for them anyway because we need them." In other words, he had made a successful business case to pay these excessive monopoly tariffs.

The fundamental question about pan-European high speed digital networks is whether the telecommunications operators will set the tariffs low enough that a large fraction of European business and industry can make a convincing business case for using them.

Putting the question of tariffs aside for a moment, why does Europe need a high speed digital network? (I use the word "network" not to describe the telecommunications operators' internal transmission networks, but a digital network service delivered directly to the customers' front door.) The basic answer is another question: why move mass from place to place when it is easier to move information? Today, the quality of telefax being what it is, we still habitually send urgent documents by courier services when perfect quality is required. But the technology exists to do much better; recently I read an announcement of a new 100-page document in my electronic mail, fetched the document from New Jersey as a *PostScript* file, and printed it (including multiple fonts, tables, and diagrams)—all within ten minutes. This is current practice for users of the Internet—the informal research network spanning the globe which has grown slowly out of the original ARPANET project funded by the US government in the 1970s.

Here today

The technology exists today to communicate electronically on a global basis between workstations, using electronic mail, electronic news systems, high-quality document transfer, data transfer, and of course all manner of remote interactive computing and transaction processing. Indeed many large international companies are using this technology on a routine basis for their internal purposes. However, a new generation of technology is coming which will move transmission speeds up by one or two orders of magnitude. This will enable a new generation of applications built around three basic capabilities: massive image transfer, instant transaction processing, and a higher quality of long distance human communication than has ever been experienced.

9

Pan-European Broadband Networks *(continued)*

Imminent network technology will allow real-time transmission of a high definition television channel between any pair of workstations; coupled with instant transaction processing this will mean that many human interactions could be carried out over long distances with almost no disadvantage compared to face-to-face contact. The inconveniences and discomforts of traditional videoconferencing will disappear.

Clearly, scientists collaborating on an international basis, such as those who use CERN, will be among the first who could benefit from such possibilities. But the same is undoubtedly true of any serious pan-European engineering project, and anyone attempting to develop a pan-European business in the open market will see the benefits.

The question is simply this: if for a given international (or even country-wide) activity, you replace the majority of human travel and physical transport of paper by use of broadband computer networking, will you get value for money?

Technology An optical fibre has a theoretically usable bandwidth of 25 *terahertz* (25 million megahertz), considerably more than the entire radio spectrum. Wavelength division multiplexing techniques, allowing the real use of a substantial part of this bandwidth, are coming out of the laboratory and have already reached the demonstration phase. And if you still happen to be short of bandwidth, a bundle of a hundred fibres can easily be made up. Furthermore, all-optical amplifiers with very high gain are now on the market. Thus, for all practical purposes, long-distance bandwidth is an infinite resource. The deployment of this bandwidth is limited by three constraints—the cost of installing fibres under the roads, the cost of transmission equipment and computer interfaces, and the cost of people to operate the network.

Bringing fibres to the premises of every business customer (not to mention domestic customers in the long run) will be an enormous investment, comparable to that made in the past when public water supply and other services were first provided. There is no way out of this if broadband services are to take off.

There is no reason to suppose that broadband transmission equipment will be overly expensive compared to earlier generations of transmission equipment. Indeed the telecommunications operators are moving to broadband technology on their trunk lines because they think it will be cheaper in the long run. There is, however, some reason for concern about the cost of computer and workstation interfaces for very high speed networks. We have become accustomed to the fact that modems or local area network interfaces for PCs or workstations are cheap, or even provided as a standard fitting. Engineering an efficient network interface to run at many Mbps is not easy and experience from the 100Mbps FDDI (*Fiber Distributed Data Interface*) standard is that interfaces cost several thousand pounds.

FDDI FDDI is the emerging high-speed LAN standard, already in service as an Ethernet backbone and now entering service as a workstation LAN. However, it is only an order of magnitude faster than Ethernet or Token Ring, and like them forces all stations to share a common slice of bandwidth. To realise the image-based computing systems of the future, a network with an aggregate capacity in the Gigabit per second (Gbps) range will be needed.

On the LAN front, the ANSI committee that defined FDDI and FDDI II is studying the requirements for an *FDDI Follow-on* (working acronym FFOL). At the same time, two Gbps channel standards—HIPPI (*High Performance Parallel Interface*) and FCS (*FiberChannel Standard*)—are being defined by other ANSI committees. HIPPI and FCS share the idea of rapid channel switching, using custom silicon to construct a switching fabric. Although both were initially conceived as computer room channels, switched HIPPI and FCS are widely considered as candidate LAN technologies. Many engineers feel that switching fabrics rather than shared media are the wave of the future for LANs.

There is a well-defined new set of standards for long-distance digital transmission, known in Europe as the *Synchronous Digital Hierarchy* (SONET in the USA). SDH defines, among other things, a set of bit rates for long distance transmission, including 150Mbps, 600Mbps, and 2.4Gbps. SDH is backwards-compatible with the older transmission standards whose "magic numbers" in Europe are 2Mbps, 34Mbps and 140Mbps. However, the SDH rates constitute a worldwide CCITT standard. The FDDI committee has decided to use SDH rates for FFOL, but unfortunately neither HIPPI nor FCS are aligned on SDH rates.

ATM The most talked-about new standard for high speed transmission is *Asynchronous Transfer Mode* (ATM), whose basic idea is to exploit the same switching fabric technology envisaged for HIPPI or FCS, but on a vastly greater scale. In ATM, data (including digitised voice) are fragmented into small packets known as "cells" which can carry up to 48 useful bytes, and routed through a network of ATM switches over pre-established virtual circuits. Because the switches are built in pure hardware, transit delay, errors, and cell losses will all be held to a very low level. On top of the ATM cell transmission service, adaptors will provide services such as telephony, 2 Mbit/s leased lines, or connectionless (LAN-style) data network service. Although ATM was invented as a WAN technology for both data and telephony usage, serious consideration is being given to applying it as a high-performance LAN.

As an intermediate step, the IEEE 802.6 *Metropolitan Area Network* (MAN) standard and the *Switched Multimegabit Data Service* (SMDS) standard, or its European variant *Connectionless Broadband Data Service* (CBDS), are being promoted. SMDS with MAN will offer a connectionless (LAN-style) service to its subscribers; SMDS islands are expected to be interconnected by ATM.

All this is not without controversy. Which Gigabit LAN will win in the market: HIPPI, FCS, FFOL, ATM or some outsider? Will all telecommunications operators offer ATM right into customer premises, or will they insist on hiding ATM behind SMDS? How will the many gaps in the new standards be filled? Will ATM appear first in the WAN or the LAN market? And will either ATM or SMDS be a commercial success, rather than becoming an embarrassment to the telecommunications industry on the same scale as narrow-band ISDN?

Almost all of these questions hinge on money. For ATM and SMDS in particular, the main issue will be whether the tariffs make them more attractive than leased lines. All past experience with switched services such as X.25 or ISDN suggests the opposite.

Pan-European Broadband Networks *(continued)*

Politics
I believe that high-speed digital networks have a lot to offer European science, industry, and business, if Europe can only provide itself with an adequate infrastructure accessible at reasonable tariffs. How can we get there from here?

The Americans believe they have solved this conundrum. Congress and the President decided that data networking and high performance computing are strategic issues for the US economy, and that the way should be shown by creating a *National Research and Education Network* (NREN). In December 1991, President Bush signed the *High Performance Computing Act,* which allocates $2.9B over five years, much of it for the NREN. Even before this, the US Government, via DARPA and NSF, had sponsored five wide-area *Gigabit Testbeds,* each involving academia, industry, and the telecommunications carriers. Thus, while European R&D struggles to obtain transcontinental 2Mbps connectivity, our American colleagues have an operational 45Mbps network with active plans for a Gbps future.

Players
Who are the European players who might help us catch up? I believe that the US model, whereby academia, and in particular scientific applications, lead the way with Government assistance, is by far the most hopeful approach. This means that we must consider the following organisations:

- The existing national and international research networks, too numerous to list here. The most notable international networks are EARN (*European Academic Research Network*) and EUNET (*European Unix Network*), both running on a not-for-profit basis. All these networks together reach at least 200,000 computers in the European research community.

- RARE (*Réseaux Associés pour la Recherche Européenne*), the international association of these networks. RARE is currently promoting the creation of an *Operational Unit* for international research networks, in the form of a not-for-profit company.

- COSINE (*Cooperation for Open Systems Interconnection Networking in Europe*), a Eureka project for X.25-based research networking which is now drawing towards its close. COSINE has created IXI (Initial X.25 Interconnect), an out-sourced private international X.25 network for R&D users.

- *EBONE 92,* a consortium of several research networks and other organisations which operates a European IP backbone at speeds up to 2048Kbps. The IP networks drawn together by EBONE can be viewed as the European component of the world-wide Internet.

- *Directorate-General XIII of the European Commission* (DG XIII). DG XIII patronised both RARE and COSINE, and more recently has supported a consultative body called ECFRN (*European Consultative Forum for Research Networking*) and the *High Performance Computing Advisory Committee* chaired by Professor Rubbia of CERN. DG XIII is also the sponsor of the RACE (*Research in Advanced Communications for Europe*) programme which has supported R&D, but not deployment, of broadband technology.

Barriers
With all these bodies in existence, what prevents us moving forward to 2Mbps and higher speeds? The barriers one can identify include:

- Confused & restrictive European telecommunications regulations
- High tariffs

- Lack of competing pan-European providers of leased lines
- General conservatism about telecommunications
- Disaccord about protocol issues (especially suitability of X.25)
- Lack of European scientific and engineering computer makers
- Lack of pan-European collaboration between users, suppliers, and telecommunications operators
- Lack of a single focal point to drive things forward

Solutions

What can be done to overcome these barriers? Suggestions include the following:

- ECFRN (see above) has proposed the creation of a *High Level Officials Group,* convened by the European Commission, to focus governmental attention on the topic.

- Debate should be shifted from political and protocol issues to the provision of high-speed open network services for the R&D community, as a stimulus to European R&D and as a pacemaker for the whole of industry and business.

- Collaborative ventures between users, suppliers and operators should be launched. Two good starts would be a 34Mbps international backbone between major research centres, and two or three international Gigabit Testbeds based on ATM.

- More liberal regulations, leading to competitive supply of international bandwidth, are a must.

- RARE, with the support of the future High Level Officials Group, should concentrate on strategy, leaving technical questions to its proposed Operational Unit. Both should concentrate on open networking, but neither should apply any form of protocol "religion."

These measures, if pursued vigorously and backed by adequate EC funding, would allow Europe to catch up with the American plans for deployment of continent-wide data networking. The following phase, extending this deployment from the research community to the whole industrial and business community, depends essentially on an economic question: will the broadband telecommunications tariffs be such as to allow managers to make an effective business case for the use of this revolutionary technology?

Acknowledgements

References

[1] "Réseaux Associés pour la Recherche Européenne (RARE)," *ConneXions,* Volume 6, No. 1, January 1992.

[2] Goldstein, S. & Michau, C., "Convergence of European and North American Research and Academic Networking," *ConneXions,* Volume 5, No. 4, April 1991.

BRIAN E. CARPENTER has been Group Leader of the Communications Systems group at CERN since 1985, following ten years' experience in software for process control systems at CERN, which was interrupted by three years teaching undergraduate computer science at Massey University in New Zealand. He holds a first degree in physics and a Ph.D. in computer science, and is an M.I.E.E. He can be reached as: `brian@dxcern.cern.ch`

When is a Network X.25?
And what is it if it isn't: a very British net?

by Jon Crowcroft, University College London

JANET

JANET, The Joint Academic Network in the UK, was a real X.25 network based on work with British Telecom on the *Experimental Packet Switched Service* which partly developed the original X.25 definition—that of a Service, *not* a Protocol.

Internally this network was provided by a number of leased lines and backbone switches. Hosts and campus switches were attached to the backbone. Most of the original network was 9.6Kbps or 64Kbps

JANET II

JANET II was firstly an upgrade to 2Mbps backbone switching and lines, and secondly a step towards more clear definition of protocols, including widespread definition and implementation of the *Coloured Book* protocols, which were a substitute for the yet-to-be-defined ISO OSI protocols, and a direct alternative to the Internet Protocols.

Internally, the backbone ran a proprietary datagram protocol between the switches and provides either level 2 (raw) HDLC, or 3 access with interface service X.25 semantics *or else* full end-to-end X.25 protocol semantics, depending on configuration of switch ports. It was one of the few X.25 networks in the world to offer actual user throughput around the several megabits per second mark. Towards the end of its life (which is not yet!), the migration to full OSI was supposed to occur as defined in the White Book (shades of Tolkien).

Instead...

JIPS and DECnet

Pressure from the plethora of systems on LANs at each campus that had DECnet and TCP/IP bundled and cost money to get Coloured Book, and serious money for toy ISO OSI implementations, caused the establishment of limited and controlled access for these protocols. The *JIPS* (JANET IP Service) and DECnet were added by building virtual nets above the X.25 "cloud" that made up the JANET II backbone.

After some performance testing, the parameters were set so that IP experienced little in the way of packet size or window size problems in traversing this net (and little in the way of extra variability in delay caused by X.25 level 3 retransmits since the net is built out of low error rate leased lines). However, it transpired that many of the routers could not run X.25 in hardware, and suffered grievous performance hits from processing the X.25 protocol as well as trying to forward IP packets. This led to the reconfiguration of some of the loaded paths to use only level 2 X.25 on those ports to which the routers were attached.

SuperJANET

SuperJANET is a twofold network with a twofold purpose. It comprises a multi-site (40 odd sites with 34Mbps) SMDS backbone to upgrade the JANET II backbone, and may be configured in a very similar way with routers around the edge of this cloud, although it is possible to have less regular topologies. It also includes a smaller number of sites (12) with 140Mbps rising to 600Mbps SDH access lines, via a star network, which in the first are "hard multiplexed," with a mix of routers and ATM switches at each site. Later, the centre may become an ATM switch, and all the end points then ATM multiplexed. This ATM "cloud" will then merge with a future evolution of SMDS to provide a seamless ATM backbone cloud.

The future

While we foresee some sites maintaining IP access, we also envisage more and more sites migrating their LANs to Hub nets or ATM LANs, and having a full end to end ATM system. This not the least so that we can support multimedia applications across the entire system without the need for ad hoc translations at the edge of the backbone.

In the interim, we will be conducting both service and research networking over this interesting mixture of technologies with the same mix of service and protocol.

Swapping layers

Meanwhile, how are the old 64Kbps and 2Mbps X.25 services provided over this high speed network? The routers provide X.25 interfaces to external hosts or switches *above* TCP connections across the backbone! A complete inversion of the original JIPS conception...

References

[1] Malis, A., "Multiprotocol Encapsulation over Frame Relay," *ConneXions,* Volume 6, No. 8, August 1992.

[2] Kozel, E., "The Cisco/DEC/NTI/StrataCom Frame Relay Specification," *ConneXions,* Volume 5, No. 3, March 1991.

[3] Vair, D., "Components of OSI: X.25—the Network, Data Link, and Physical Layers of the OSI Reference Model," *ConneXions,* Volume 4, No. 12, December 1990.

[4] Falk, G., "LAN Interconnection Over X.25," *ConneXions,* Volume 7, No. 2, February 1993.

[5] Hughes, L., & Starliper, S., "Switched Multimegabit Data Service (SMDS), *ConneXions,* Volume 4, No. 10, October 1990.

[7] Krishnaswamy, P., & Ulema, M., "Developments in SMDS," *ConneXions,* Volume 5, No. 10, October 1991.

[8] CCITT Temporary Document 39, Melbourne, Australia, December 1991, WPXVIII/8 Meeting Report—Annex 3 (Part 2), "Text of WP XVIII/8 B-ISDN Recommendations," Appendix 5, Draft Recommendation I.150, "B-ISDN ATM Functional Characteristics."

[9] CCITT Temporary Document 39, Melbourne, Australia, December 1991, WPXVIII/8 Meeting Report—Annex 3 (Part 2), "Text of WP XVIII/8 B-ISDN Recommendations," Appendices 9 & 10, Draft Recommendation I.363, "B-ISDN ATM Adaptation Layer (AAL) Specification."

[10] ANSI T1S1.5 / 91-449, "AAL5—A New High Speed Data Transfer AAL," November 1991.

[11] IEEE Standard 802.6-1990, "Distributed Queue Dual Bus (DQDB) Subnetwork of a Metropolitan Area Network (MAN)."

[12] Bellcore Technical Reference TR-TSV-000772, Issue 1, May 1991, "Generic System Requirements in Support of Switched Multi-Megabit Data Service."

[13] Clapp, G. and Zeug, M., "Components of OSI: Asynchronous Transfer Mode (ATM) and ATM Adaptation Layers," *ConneXions,* Volume 6, No. 4, April 1992.

[14] Crowcroft, J., "UK Internet Protocol Connectivity Activities," *ConneXions,* Volume 5, No. 6, June 1991.

JON CROWCROFT is a Senior Lecturer at University College London where he has been engaged in Internet related research for about 10 years. He got his BA from Cambridge some time ago, his Masters from London more recently, and intends completing his PhD before Cyberspace is fully colonized. He can be reached as: jon@cs.ucl.ac.uk

You cannot promote OSI Applications over OSI Networks

by Paul Barker, University College London
and
Colin J. Robbins, NeXor Ltd.

Introduction

The authors of this article have both spent several years developing OSI applications. During the course of this period we have attempted to demonstrate our software at a number of exhibitions, where a prime goal was to convince the attendees that OSI applications were now available and usable. Due to the politics of OSI being as they are, we have almost always been obliged to demonstrate our OSI applications over a full OSI protocol stack; the lack of a *Connection-Oriented Network Service* (CONS) or *Connectionless Network Service* (CLNS) network infrastructure has meant that X.25 (1980) has been mandated as the network service.

This article argues, based on a catalogue of frustrating, not to say embarrassing, experiences that the insistence of using X.25 as a network service is positively injurious to the acceptance of OSI applications. The evidence presented here suggests that X.25 implementations are of poor quality when compared with other networking technology. We have encountered problems with lack of robustness, configuration difficulties and poor support. Furthermore, it is often not straightforward to physically connect a machine to an X.25 network; again, an indication of the problems will be given. We also examine whether the migration to CONS is likely to improve matters.

To set the problems in context, we describe an experience at another exhibition, to show how easy it ought to be. Finally, we draw conclusions as to why things are the way they are and note the steps needed to remedy the situation. We suggest that the cause of OSI may be better served by, for the moment at least, using non-OSI networking technology.

Notes: throughout this article we refer to X.25. This is a reference to X.25 (80), unless stated otherwise, as this version of the protocol is the one used most often. The experiences we report have all been with X.25 implementations on UNIX machines. This article contains the personal opinions of the authors. It does not reflect the policy or opinions of their respective organisations.

Don't forget your toolbox

When attempting to use X.25, the first thing to be considered is connecting your machine (DTE) to the X.25 network provided at the site (DCE). It is rarely a simple matter of just plugging in and away you go. In fact, you probably won't get very far without this lot: soldering iron; pliers; screwdriver; breakout box (with spare battery); universal interface converter; protocol analyser; bag of plugs, cables and adaptors.

There are a number of reasons why such a formidable tool-kit is essential. First, the gender of connectors is not consistent across different makes of hardware. Whilst the use of male for DTE and female for DCE is common, it is far from ubiquitous.

For reasons which are beyond the authors, but seem to be accepted as normal by technicians, looping pairs of pins together seems the norm rather than the exception when using V24 connectors. The breakout box is an essential tool here in diagnosing which loops and crossovers are required in any instance. Of course, breakout box batteries don't last forever and so don't forget to take a soldering iron to produce a finished cable.

There are a variety of different connector types with which you could be presented. The type of connector differs depending on the offered line speed:

- RS 232 / V24 /X.21bis: 25 pin connector — up to 19.2Kbps
- V11 / X.21 Interface: 15 pin connector — up to 64Kbps
- V35 / X.21bis interface: 34 pin connector — up to 48Kbps
- RS 449: 37 pins — up to 2Mbps

Unless you are absolutely sure of the equipment to which you are connecting, and even, to be prudent, if you are sure, the only safe way round this maze is to take along a universal interface converter. These magic boxes are not cheap, but at least you should be able to connect your machine to the network.

Even when you get a cable that works, your problems may not all be over. We have found that an X.25 protocol analyser is very helpful to sort out what is going wrong.

Problems with X.25 implementations

The following sections describe some of the problems encountered trying to set up an X.25 link due to limitations of various manufacturers' implementations.

Clock synchronisation

Both the DCE and DTE should be able to supply the clock signals required by the connection. However, one manufacturer insists that their machine acting as a DTE supplies the clock signal. But, many modems or DCEs also expect to supply the clock. Consequently this manufacturer's machines cannot be connected to many DCEs. (We have subsequently been told that theoretically this shouldn't be the case. But, since not all pins are always connected, and software doesn't always look at every signal, it's very hard to know whether we could have got it working properly if we'd had the time to persevere.) In fact, we did get something working. We had a 9.6Kbps X.25 line. Both DTE and DCE supplied clocks and were not synchronised. However, every now and again the clocks would run together, and we managed to get some packets through. Since we were running TCP/IP over this link, with its error recovery facilities, it was able to cope with the 90% packet loss. We were effectively left with a 1Kbps line for the demonstration—but better than nothing!

Addressing

Addressing application entities is a fundamental part of any networking system. X.25 addressing is based upon DTEs or X.121 addresses, but the addressing does not work consistently across international boundaries. For example UCL's public network (PSS) DTE address is:

234219200300

This contains a DNIC or country code of "234," which identifies the UK. From the UK if you call this address, you will get a connection. Similarly, the PSS address of GMD in Germany is:

262450502303

In this case the DNIC is "262." If you call this from the UK you will get a connection. However, if you are in Germany and wish to call these services, you have to modify these numbers. For the UK address (and all international calls) you have to add a leading zero:

0234219200300

and for the local address, the DNIC has to be stripped:

450502303

17

OSI Applications over OSI Networks *(continued)*

If these addresses may be configured in a local table, there is no problem. However, if the addresses are obtained from a directory service such as X.500 where only one form of address is stored, the implication is that software has to know about these local addressing differences and add the zeros, or strip DNICs as appropriate. This makes the configuration steps complex.

The leading zero also causes problems with private networks such as the European IXI network. For IXI the pseudo DNIC "204" is used. In most countries you treat IXI as an international call, and add the leading zero if appropriate. However this is not always the case. In Switzerland, for example, all international addresses need a leading zero, but *not* IXI addresses. This means software has not only to recognise its local DNIC, but also some other DNICs as pseudo local. Not all X.25 implementations can cope with this.

Configuration

One major problem is configuration of X.25 parameters. Not only is it complex, but implementations often lack sufficient flexibility. The address format problem described above is a typical case. We are familiar with some application software that knows about these problems and has simple tailoring to allow the options to be specified. However, the interface provided to the operating system does not always allow this information to be passed through the API, with the operating system having its own tailoring for these parameters. This adds confusion and frustration to people attempting to configure such a system.

Sub-addresses

X.25 sub-addresses cause confusion. If a sub-address is used, some X.25 switches can remove the local DTE from the full address when passing to an application, leaving just the sub-address. Unless the program invoking the network listen is aware of this and just listens on the sub-address, the connection will not be made.

Logical Channels

Logical channels are probably the black art of X.25. Unless the DTE and DCE agree about which logical channels can be used for the connection, nothing works. This would be acceptable to some extent if it was possible to give error information of the "invalid channel" form. However, our experience is connections just hanging, with no indication of the problem. This just adds to the complexity of setting up a connection.

We have also seen misunderstanding over the use of logical channel zero. (Annex A of the X.25 standard says that logical channel zero can not be used by an SVC, but can be used by a PVC.) At a recent demonstration, we attempted to connect a machine to a PSS line provided by that country's PTT. The PSS DCE used supplied logical channels 0 to 3 (for an SVC). However, logical channel 0 was not supported on the machine we had. This meant the first connection attempt always failed, and a simple program had to be quickly written to use up the first channel to get the demonstration to work at all.

Window size

The window size has caused some baffling problems. A recent example of this was when we tried to access a UK service from Belgium. We had a window size of 2 configured. When we tried to make a connection, some protocol was passed, but then the connection hung. This was traced at the remote end to the "more" bit being set, but no more data being sent. Adjusting the window size to 7, cured this. Apparently, the Belgian PTT could only accept a window size of 7. X.25 has the protocol support to negotiate these parameters, but it is not mandatory, and so rarely implemented.

Fast Select

Fast Select is an optional part of X.25. However, we have seen an OSI stack implementation which assumes fast select is available. However, unless the remote site supports this option it cannot be used. Furthermore, not all PTTs support fast select even if the two machines attempting to communicate do.

Handling software interrupts

One machine, of which we have had experience, has problems with software interrupts; usually these are issued in an attempt to regain control of a process which is hanging. In one instance, if we interrupted a UNIX process making an X.25 call by issuing a control-C, the X.25 driver was left in an odd state. The next time an X.25 connection was attempted, the machine panicked and rebooted. We should consider ourselves lucky. Some colleagues have reported to us that their machine will not reboot in the same circumstances unless the X.25 card is first removed! It is hard to have much faith in the overall system when confronted with such problems.

Integration with O/S

A number of problems stem from the fact that X.25 is often not provided as part of the operating system software, whereas TCP/IP, for example, is fully integrated. There are two elements to this separation: the psychological; the technical. Psychologically, the provision of X.25 outside of the mainstream operating system release carries a penalty in terms of software quality and support. However, a more serious problem is that sometimes the software is not technically aligned either. A standard facility for handling asynchronous network events on UNIX is the *select()* system call. So long as this interface is used, it is a simple matter to manage listeners on a number of networks. However, if an X.25 implementation uses a different interface, it becomes much more difficult to manage asynchronous network events. Not only does this complicate porting, but there may also be a performance penalty as the listening process has to alternate between the different styles of network listener.

A protocol problem

We have cited a large number of examples highlighting basic inadequacies in manufacturers' X.25 software, reflecting the fact that X.25 is not a "flagship" product. It is worth briefly noting a protocol problem which has caused us many headaches in the past. If you attempt to start a network listen program using TCP/IP services, and you get the address wrong, you get an error of the form:

```
Can't assign requested address
```

Furthermore, if the address is correct, but in use by another listener you get:

```
Network address in use
```

These are two very important error conditions. The first message indicates that a configuration detail needs modifying; the second that the system has detected a condition which could cause unpredictable behaviour. In X.25 it is not possible to detect either of these conditions. This is considered by the authors to be a major design flaw in the protocol.

Support

The quality of the implementation is, in a sense, just a starting point: products with problems can be redeemed by energetic support; good quality products can similarly be ruined by lack of support. Our experience with several manufacturers is that X.25 support is inadequate. Often, the only way to make any progress is to get a contact on their development team but, understandably, this is not always easy to do.

19

OSI Applications over OSI Networks *(continued)*

During one recent demonstration we needed to contact a large multi-national corporation for assistance. We had a contact, but he was on honeymoon. Through another, non-technical, contact we discovered that this corporation had run into a similar problem the week before whilst trying to do a similar demonstration. They, too, were awaiting the return of the honeymooner! Is it credible that their mainstream networking software would be so poorly supported?

We have also found it hard to obtain the appropriate level of support attention. We once reported a bug in an X.25 implementation, which the manufacturer eventually agreed was a problem. (This was a feat in itself, and required the writing of a test program to demonstrate the problem.) A week later we got back the request "Can you change line X of your program to... and try again?" This response hardly inspired confidence and, predictably, the change did not solve the problem. A few months later, we did receive a fix, but heard from an inside source that a large American PTT just reported the same problem—the fix appeared in less than 24 hours. Our cynical interpretation is that X.25 often gets something akin to "best effort" support, unless the customer has market clout. This fits in with the feeling that X.25 is regarded as peripheral to manufacturers' business.

One manufacturer is currently encouraging their customers to migrate to a new operating system. New equipment is delivered with the new operating system software. However, their X.25 software is not available for the new operating system, and will not be for several months.

CONS The problems discussed so far relate to X.25 (80). Some countries are now beginning to run X.25 (84) services, or CONS (Connection Oriented Network Service). This provides a whole new set of problems, both in terms of implementations and protocol. We have offered much anecdotal evidence of the state of X.25 implementations, and only say here that on current evidence that the situation is as bad, if not worse, for CONS. Most of the reservations which people have over CONS are fundamental addressing issues. The main problems are as follows:

Addressing is based upon NSAPs, which can map onto DTE addresses, or more generically SNPAs. There is no mechanism defined for providing this mapping. Each manufacturer is providing their own proprietary solution, if they have a solution at all. Ad hoc solutions based on tables, or centralised databases such as the UK academic community's *Name Registration Scheme,* have inherent problems of scale.

OSI does not describe one network services, but two: CONS and CLNS (Connectionless Network Service). However, both of these use the same NSAP address space. When an application is presented with an NSAP, there is no way, without using external knowledge, of determining if it represents a CONS or CLNS based service. The best you can do at present is to try and access the service over either CONS or CLNS, and see if it works. Of course, this is unbelievably crude, and colours many people's view of OSI in general.

How it should work Connecting a machine to a network really need not be so difficult. Consider the following example which, as for all the previous instances, is based on experience:

We turned up at a conference with a standard machine, in which the network interface was a standard part of the operating system. We took two cables, knowing one would fit for sure. We plugged the system together, changed the address of the machine in a single file and restarted the network daemon. We then ran our application, and it all worked straight away. We had just set up a TCP/IP link from a UNIX machine to an Ethernet, which was linked to the global Internet.

Conclusions

We are the first to admit that we are not X.25 experts. We are OSI application developers, wishing to use OSI networks for our applications. With the current state of X.25 technology we have had to acquire a large amount of practical X.25 knowledge to get these applications to work. This need not be the case. We feel confident about being able to connect a machine to a piece of Ethernet and configure TCP/IP to run over it, then being able to configure our OSI application to run over TCP/IP using RFC 1006. This is *not* because we are TCP/ IP experts (probably the opposite is true—we now know more about X.25 than TCP/IP), but because TCP/IP is fully tried, tested and supported by manufacturers, and is a fundamental part of many operating systems.

It should be no mystery to anyone why the situation is the way it is. Low quality X.25 implementations are symptomatic of computer (software) manufacturers following a different market trend. It is unrealistic to expect them to invest heavily in X.25 software development, whilst their customers have largely decided to "watch this space" as far as OSI networking is concerned, and decided to use TCP/IP, or some other well-supported proprietary technology.

The motivation for writing this is that we regard the mandating of X.25 (or CONS or CLNS for that matter if they too are not taken seriously by the market) for providing the network service as detrimental to the cause of OSI in general. There is a very real danger that aspects of OSI which currently have a good measure of support will, to use the vernacular, be thrown out with the OSI networking bath-water.

There is an alternative, and this idea has to be accepted by those who have influence over networking policy. Until OSI networking comes of age, it is more beneficial to the overall cause of OSI to allow OSI applications to be run over widely used, reliable networking technology. Techniques, such as the one described in RFC 1006, allow the use of OSI applications over non-OSI layers. Such an approach is manifestly successful; much of the current usage of OSI applications, away from the exhibition halls, is based on non-OSI networks. It is ironic that OSI politics hinder us in trying to demonstrate our working OSI applications at precisely those exhibitions which seek to promote OSI. We regard the following as necessary conditions for the long-term viability of OSI networking, and by implication, OSI in general:

- Implementors' agreements and profiles on options and parameters;

- X.25 as a turn-key service, as an integral part of the operating system;

- X.25 software configuration must be easier;

- High quality support;

- Resolution of CONS/CLNS addressing problems.

21

OSI Applications over OSI Networks *(continued)*

It is possible that these steps will be taken, although it seems unlikely. Coercion by governments, whereby contracts are only offered to those manufacturers able to offer OSI networking, has resulted in nominal, poor quality implementations, which meet contractual requirements alone. Ultimately, only the market can save OSI networking, and currently the market looks very sceptical. Our view is that the OSI applications will also fail to find favour so long as they are inextricably associated with OSI networking. Until OSI networks offer the same ease of use as existing TCP/IP and proprietary networks, the overall cause of OSI is better served by allowing OSI application developers to use the best networking technology available.

References

[1] Vair, D., "Components of OSI: X.25—the Network, Data Link, and Physical Layers of the OSI Reference Model," *ConneXions*, Volume 4, No. 12, December 1990.

[2] Falk, G., "LAN Interconnection Over X.25," *ConneXions*, Volume 7, No. 2, February 1993.

[3] "Information processing systems—Telecommunications and Information Exchange between systems—Protocol for Providing the Connectionless-mode Network Service," ISO 8473, March 1987.

[4] Marshall T. Rose and Dwight E. Cass, "ISO Transport Services on top of the TCP," RFC 1006, May 1987.

[5] Hagens, Rob, "Components of OSI: CLNP or A Day in the life of Ivan CLNPacket," *ConneXions*, Volume 3, No. 10, October 1989.

[6] Callon, R. "An Overview of OSI NSAP Addressing in the Internet," *ConneXions*, Volume 5, No. 12, December 1991.

[7] Malis, A., "Multiprotocol Encapsulation over Frame Relay," *ConneXions*, Volume 6, No. 8, August 1992.

[8] Kozel, E., "The Cisco/DEC/NTI/StrataCom Frame Relay Specification," *ConneXions*, Volume 5, No. 3, March 1991.

[9] Krishnaswamy, P., & Ulema, M., "Developments in SMDS," *ConneXions*, Volume 5, No. 10, October 1991.

[10] Lippis, N., & Herman, J., "Widening Your Internet Horizons: Wide-Area Options for Internets," *ConneXions*, Volume 5, No. 10, October 1991.

[11] Lippis, N., Moore, S., & Morency, J. P., "Multiprotocol Internets: Users' Best Hope for Open Systems," *ConneXions*, Volume 7, No. 3, March 1993.

COLIN ROBBINS received a BSc in Computer Science and Electronic Engineering from University College London. Following this he spent 3 years at UCL, where he was the primary implementor of the QUIPU X.500 system, which is widely used in both US and Europe. He is now responsible for the development of Directory and related OSI products at NeXor Ltd (formerly X-Tel Services Ltd) and is also active as a consultant with several European Commission funded OSI research and development projects. He can be reached as: `c.robbins@nexor.co.uk`

PAUL BARKER received a B.A. in Economics from Exeter University and a Diploma in Computer Science from Birkbeck College, London. He has worked in the department of Computer Science at University College in computing since 1986. During that period he has worked on a variety of projects concerned with X.500 Directory Services, and is currently investigating the use of the Directory for bibliographic purposes. He is past secretary to the U.K. Academic Directory Group. He can be reached as: `p.barker@cs.ucl.ac.uk`

MICE

Objectives

Esprit Project 7602, *Multimedia International Conferencing for Euro-pean Researchers* (MICE), commenced work on 1 December 1992. The major objectives of the project are:

- To provide appropriate multi-media, multi-party conferencing facilities to research workers in Europe, with links to North America; these to be provided across disciplines and ideally, usable by any research worker.

- To design these facilities to interwork with those already deployed for other projects in the research community.

- To provide the facilities in a form that would be appropriate to complement multi-party meetings so that one emphasis will be on conference rooms; workstations with reduced facilities will also be embraced. In both cases the facilities provided will include shared workspace use of workstations as well as video-conferencing.

- To design the facilities to be deployable with minimal expenditure on communications and minimal delay, thus precluding the provision of a separate communications network.

- To use the emerging European infrastructure of international connections between National Packet Switched Networks (Europanet and EBONE inside Europe, and the different links to the US from Europe).

Implementation

The project is planned for completion in 12 months. It will consist of three overlapping phases: definition, trials and evaluation. During the definition phase, a multimedia conferencing reference architecture will be defined; thereafter the exact facilities to be provided in conferencing rooms, conferencing workstations and a Conference Multiplexing and Management Centre will be specified.

During the trials phase, the facilities of all three areas will be improved progressively and put into limited service. During the evaluation phase, the strengths and weaknesses of the facilities will be assessed, forecasts will be made on the timing of more economic equipment and more appropriate algorithms, and recommendations will be made for the deployment of an operational system. The timescale is only achievable, within the limited effort proposed, because the project will be based heavily on existing developments funded under other programmes.

Participants

The MICE consortium includes the following partners: University College London (UCL) (prime contractor), Stuttgart Univeristy, INRIA, University of Oslo, ONERA, Swedish Institute for Computer Science, GMD, Nottingham University, ULB/VUB, NTR.

It is intended that the initial four sites should be connected within the first four months of the project.

More information

For more information about MICE, please send e-mail to:

```
c.easteal@cs.ucl.ac.uk      ; Management

a.sasse@cs.ucl.ac.uk        ; Human Factors

p.kirstein@cs.ucl.ac.uk     ; Policy issues

m.handley@cs.ucl.ac.uk
           and              ; Technical issues
s.clayman@cs.ucl.ac.uk
```

23

Announcement and Call for Participation

The 7th *USENIX Systems Administration Conference* (LISA VII) will be held November 1–5, 1993 at the Marriott Hotel in Monterey, California. The annual LISA conference provides a forum in which system administrators meet to share ideas and experiences. A growing success for the past six years, LISA is the only conference which focuses specifically on the needs of system administrators. Its scope includes system administrators from sites of all sizes and configurations.

Tutorial program

The two-day tutorial program (Monday and Tuesday) at LISA VII is divided into three tracks with a total of twelve half-day tutorials. Attendees may move between tracks, choosing which sections most interest them. Tutorials offer expert instruction in areas of interest to system administrators, novice through experienced. Topics are expected to include Networking, Programming in *Perl,* The X Window System and the Administrator, the Domain Name System, *Sendmail,* and more.

Technical sessions

"The Human Aspect of UNIX System Administration" is the theme of the first track of the conference technical sessions. Although papers of a more traditional technical content are also very welcome, the committee is especially seeking papers on areas such as creating policies and procedures, interacting with management and/or users, or which describe and evaluate methods aimed at improved communication with users and/or management. We are interested in papers which provide freely available or fully described solutions to existing problems.

The second track of the conference technical sessions will be split between presentations on very large installation system administration and presentations of practical, experience-derived material of special interest to new system administrators.

No simple measure defines "large installation." It could be number of hosts, number of users, size of network, amount of on-line storage, or some combination of these. The only certainty is that today's "large" will be tomorrow's "standard." We wish to hear from sites which have unique problems and solutions relating to the management of large installations. We seek proposals for papers, panels, mini-workshops, or similar presentations for this track.

We also seek papers, mini-workshops, or panel presentations of pragmatic material from experienced system administrators who wish to share their experiences, hardships and solutions. It is hoped that administrators at all levels can leverage this track to solve specific problems at their site.

Vendor display

Well informed vendor representatives will demonstrate products and services useful to system and network administration at the informal table-top display accompanying the LISA Conference. The display will be open Wednesday, November 3, 1993, 5 pm – 9 pm. If your company would like to participate, please contact Cynthia Deno at telephone 408-335-9445, FAX 408-335-2163, E-mail: cynthia@usenix.org

Topics

The technical sessions program may include invited talks, panels, *Works-In-Progress* (WIP) reports, and *Birds-Of-a-Feather* (BOF) sessions, alongside the refereed paper presentations. The program committee invites you to submit informal proposals, ideas, or suggestions on any of the following or related topics:

- Implementation, usage, & strategies for policies and procedures
- Effects of improved communication with users and management
- Tricks in user education
- How to develop Junior System Administrators
- System Security Monitoring
- Security issues, where multiple people are privileged users
- Heterogeneous system administration
- Human issues of administration
- Graphical User Interfaces for system administration
- Distributed system administration
- Network growth and performance management
- Network management
- Network monitoring
- Wireless LANs
- Evaluating performance of high-end workstations and servers
- Integration of heterogeneous systems
- Usage monitoring and accounting systems
- Administration of remote sites

Submissions

The committee requires that an extended abstract be submitted for the paper selection process. (Full-papers are not acceptable for this stage; if you send a full paper, you must also include an extended abstract for evaluation.) Your extended abstract should consist of a traditional abstract which summarizes the content/ideas of the entire paper, followed by a skeletal outline of the full paper. We require electronic (*nroff/troff, TeX* or ASCII) submission of the extended abstract.

Authors of an accepted paper will present their paper at LISA VII and provide a final paper for publication in the Conference Proceedings. Final papers are limited to 20 pages, including diagrams, figures and appendix. Papers should include a brief description of the site (if applicable), an outline of the problem and issues, and details of the solution. Authors may provide electronic versions or camera-ready copy (instructions will be provided upon request) of final papers. Conference proceedings will be distributed to all conference attendees and will also be available from the USENIX Association after the conference.

For submission of all proposals and of extended abstracts of refereed papers, and for program information, contact:

Bjorn Satdeva, Program Chair
2787 Moorpark Avenue
San Jose, CA 95128
Phone: 408-241-3111
E-mail: bjorn@sysadmin.com

Important dates

Extended Abstract Submission Deadline: July 2, 1993

Notification to Authors: July 23, 1993

Final Papers Receipt Deadline: September 7, 1993

Registration information

Materials containing all details of the symposium program, registration fees and forms, and hotel discount and reservation information will be mailed and posted to the net beginning August 1993. If you wish to receive registration materials, please contact:

USENIX Conference Office
22672 Lambert Street, Suite 613
El Toro, CA 92630 USA
Phone: 714-588-8649 • Fax: 714-588-9706
E-mail: conference@usenix.org

Call for Papers

The *USENIX Winter 1994 Technical Conference* will be held January 17–21, 1994 in San Francisco, California.

Introduction

For many years, UNIX and its derivatives have been the only widely available "open" operating systems to support modern software technology. The next few years promise to change that, as many new PC operating systems reach the market. These systems will compete with UNIX, but they will also broaden the set of systems that can support advanced applications, high-performance computing, novel user interfaces, and improved network communication. The question is not "will UNIX survive," but rather how will UNIX and other systems evolve together to improve our computing environments.

As usual at USENIX Conferences, we are interested in papers describing new and interesting developments in open operating systems. Our traditional focus on UNIX remains, but this includes lessons learned from work on UNIX that can be applied more broadly, and lessons from other kinds of systems that can be applied to the continuing evolution of UNIX.

Important dates

Extended Abstracts Due:	July 13, 1993
Notification to Authors:	August 30, 1993
Camera-ready Papers Due:	November 2, 1993

Instructions for authors

As at all USENIX conferences, papers that analyze problem areas and draw important conclusions from practical experience are welcome. Note that the USENIX conference, like most conferences and journals, considers it unethical to submit the same paper simultaneously to more than one conference or publication or to submit a paper that has been or will be published elsewhere.

Cash prizes will be awarded by the program committee for best paper, best presentation and best student paper.

Authors of papers to be presented during the conference technical sessions and published in the Proceedings must submit an extended abstract to the Program Committee by July 13, 1993. The object of an extended abstract is to convince the reviewers that a good paper and 25-minute presentation will result. The reviewers need to know that authors:

- Are attacking a significant problem;

- Are familiar with the current literature about the problem.

- Have devised an original solution;

- Have implemented it and, if appropriate, characterized its performance;

- Have drawn appropriate conclusions about what they have learned and why it is important.

An extended abstract should be about 5 pages in length, or about 2500 words. (Final papers should be 8 to 12 pages long.) The extended abstract should represent the paper in "short form." It should include the abstract as it will appear in the final paper. The body of the extended abstract should be complete paragraphs, not just an outline of the paper. (Sections present in the full paper but omitted from the abstract may be summarized in terse form; this will help reviewers to understand what material will be present in the final paper.)

Authors should include full references, and, if appropriate, performance data to establish that they have a working implementation and measurement tools. Figures should be included when available.

Authors may, at their option, submit a full paper *in addition* to the extended abstract. Since the schedule for reviewing submissions is short, however, reviewers do not have time to read full papers for all submissions, and most judgments will be made based on the extended abstracts.

Every submission should include one additional page containing:

- The name, surface mail address, daytime and evening telephone numbers, e-mail address and (if available) fax number of one of the authors, who will act as the contact to the program committee;

- An indication of which, if any, of the authors are full-time students;

- A list of audio/visual equipment desired beyond a microphone and an overhead projector.

Authors of accepted submissions will be notified by August 30, 1993. They will promptly receive instructions for preparing camera-ready copy of an 8-12 page final paper, which must be received by November 2, 1993.

Where to send submissions

Please submit one copy of an extended abstract via at least two of the following methods:

- (Preferred method) e-mail to: SF94papers@usenix.org

- Fax to: the USENIX Association at +1 508-548-5738

- Mail to:

  Winter 94 USENIX
  USENIX Association
  2560 Ninth St., Suite 215
  Berkeley, CA 94710
  USA

Inquiries about submissions to the USENIX Winter 1994 Conference may be made by e-mail to: SF94info@usenix.org or by telephone to +1 510-528-8649. Potential authors of technical papers are strongly encouraged to send us electronic mail. This will allow us to notify you of any important changes and you will receive additional information about the submission and reviewing process.

You may request a sample extended abstract by telephone to +1 510- 528-8649, by fax to +1 510-548-5738, or by e-mail to: sample-abstract@usenix.org

Conference program and registration

Materials containing all details of the technical sessions and tutorial program, conference registration, hotel discounts, and airfare discount and reservation information will be mailed at the end of September 1993. If you wish to receive the registration materials, please contact:

    USENIX Conference Office
    22672 Lambert St., Suite 613
    Lake Forest, CA 92630
    USA
    Phone: +1 714-588-8649
    FAX: +1 714-588-9706
    E-mail: conference@usenix.org

Call for Papers

The 1994 *IFIP International Working Conference on Upper Layer Protocols, Architectures and Applications* (ULPAA) will be held May 30 – June 3, 1994 in Barcelona, Spain. The conference is sponsored by IFIP Working Group 6.5 and will be hosted by Universitat Politecnica de Catalunya.

Overview

Increasingly, the effective use of computers depends upon smooth interworking between disparate and diverse computers, communicating with each other using a wide variety of networking technologies. Such smooth interworking, in turn, depends critically on standardized communication protocols, which are themselves dependent on clearly-understood and well-specified architectures for distributed applications. As new applications are developed, new protocols are vital to the success of the applications in the world. The ULPAA conference provides a pre-standards forum where leading researchers can discuss promising and problematic developments in the world of distributed applications. Past conferences in this series have had significant impact on the earliest stages of standard development in both the ISO and Internet protocol suites. We are soliciting papers that will help to focus the efforts of the networking community and to point out new direction for continued progress.

Format

The purpose of the conference is to provide an international forum for the exchange of information on the technical, economic, and social impacts and experiences with upper layer protocols, architectures and distributed applications. The conference format will be two and a half days of conference paper presentations combined with one half a day of workshops.

Topics

Papers are desired in—but not restricted to—the following topic areas:

Application architectures:
- Implementation, and experience with distributed applications.
- Models and designs.
- Programming environments.
- Group communication models and services.
- The impact of human factors on upper layer protocol design and implementation.
- Multimedia applications and communications.
- Management and operation of distributed services.

Impact on applications of underlying services:
- Interconnection of upper layer and application entities.
- Mobile communications.
- Upper layer network management and naming.
- Presentation and session layer issues.
- Security and privacy provision.
- Very high-speed networking.

Standards:
- Upper layer conformance and interoperability testing activities.
- The role of the standardization process for upper layer protocols.

Instructions for authors Prospective authors are invited to submit for review, unpublished original contributions (not exceeding 5000 words) which describe recent research results or developments directly relevant to upper layer protocols, architectures or distributed applications.

Publication Papers that are accepted will be published by North-Holland. A preprint of the proceedings will be provided to attendees.

Important dates

Today: Send a message, letter, phone, or fax to either of the contacts below stating your intention to submit a paper, or stating your general interest in the conference.

November 1, 1993: Full version of papers due for review.

February 1, 1994: Notification of acceptance/rejection.

April 1, 1994: Camera-ready papers due for publication.

Please submit five copies of your paper to either of the Program Committee Co-Chairs:

Dr. Nathaniel S. Borenstein (North American Co-chair)
Room MRE 2D 296
Bellcore
Morristown, NJ 07962-1910
USA
Telephone: +1 201-829-4270
Facsimile: +1 201-829 5963
E-mail: `nsb@thumper.bellcore.com`

or to:

Dr. Manuel Medina (European Co-chair)
Universitat Politecnica de Catalunya
Gran Capitan, s/n. – D6.207 – Dept. Arquit. Comput.
ES-08071 Barcelona
SPAIN
Telephone: +34 3 401-6984
Facsimile: +34 3 401-7055
E-mail: `medina@ac.upc.es`
X.400:
`S=Medina/OU=ac/O=upc/PRMD=iris/ADMD=mensatex/C=es`

Tutorials On Monday and Tuesday, May 30–31, time and space is reserved for a number of tutorials. The following topics are being considered:

- Upper Layer/Application Layer Architecture
- Security Models, Mechanisms and Systems
- Message Handling X.400 (1992)
- Directory Services X.500 (1992)
- Network Management
- ASN.1
- Coexistence and Transition to OSI Applications
- Distributed Application Programming Environments
- Privacy Enhanced Messaging (PEM)
- Multimedia Internet Mail (MIME)

Call for Papers

The *Fourth International Workshop on Network and Operating System Support for Digital Audio and Video* will be held November 3–5, 1993, at the Lancaster House Hotel, Lancaster, U.K. The workshop is presented in co-operation with ACM SIGCOMM, ACM SIGOPS, ACM SIGOIS, IEEE Communications Society and IEE.

Topics

Important topics for the workshop include (but are not limited to):

- Broadband/ATM networks
- Multimedia network interfaces
- Communication protocols for multimedia
- Micro-kernel and OS support
- Application of real-time techniques
- Media synchronisation
- Quality of service support
- Multimedia storage and I/O architectures
- Distributed multimedia systems
- Integrative standards, e.g., TINA and ODP
- Performance studies

Workshop theme

Network and operating system support for digital audio and video is currently an area of intense activity, and a number of core techniques are beginning to emerge. The emphasis of this workshop will be on integration of key technologies to produce complete solutions. The role of the operating system is seen as particularly important in this respect. Papers are also welcomed on practical experiences of developing multimedia systems.

The workshop is intended to bring together practitioners from a variety of areas, including communications and networks, operating systems, real-time systems and distributed computing. It is intended that the workshop will produce an agreed position statement on the state of the art in the field, highlighting major areas requiring future research.

Instructions for authors

Authors are requested to submit a 500–2000 word position paper or an extended abstract of a full paper (in raw, unformatted text) by electronic mail to: `av-workshop@comp.lancs.ac.uk`. Only if electronic submission is impossible, papers may be sent to:

Prof. W.D. Shepherd
Computing Department, Engineering Building
Lancaster University
Lancaster LA1 4YR
ENGLAND
Fax: +44 524 38 1707
Phone: +44 524 59 3827
E-mail: `doug@comp.lancs.ac.uk`

The proceedings of the workshop will be published by Springer-Verlag and the best papers will be forwarded to selected journals for publication.

Important dates

Abstracts due:	July 19, 1993
Acceptance notification:	August 30, 1993
Final papers due:	October 4, 1993

Book Review

TCP/IP: Architectures, Protocols, and Implementation by Sidnie Feit, McGraw-Hill, ISBN 0-07-020346-6, 1993.

As the Internet has expanded, the supply of TCP/IP books has expanded with it. Not many of the new books are worth buying. But this book is a nice exception to that rule—a carefully written book on TCP/IP that tries to explain the entire protocol suite.

Organization

The outline of the book is fairly comprehensive. It starts with an introduction and history and proceeds through physical and link layers, IP addresses and Domain Name System (DNS) names, to the various protocols, IP, TCP, Telnet, FTP, NFS (with SUN RPC), and finishes with a brief primer on network administration and how to program with sockets. Overall, each topic is treated fully and professionally. I found the book generally easy to read.

Errors

The book does have a lot of errors, as any first edition survey typically does. A number of the errors are important. Karn's algorithm is completely mis-described—its purpose is get good Round-Trip Time samples for Jacobson's algorithm, not to respond to congestion. The use of DNS MX Resource Records (RRs) is incomplete—beyond indicating mail servers, MX RRs are also used for e-mail gatewaying.

The description of layering on page 26 is a canonical example of what not to say about layering—its primary utility is not as an implementation technique (as the text suggests) but rather a method for dividing up the work of standards groups. (If you don't believe me, go read Zimmerman's original paper on OSI layering [1].) And the TCP server example on pages 395–397 is unfortunate because it shows a server binding to an arbitrary port, while servers generally have to bind to well-known ports.

Useful reference

Errors aside, this is a solid book. And as it is refined and improved in later editions, I would expect it to become a useful reference on one's bookshelf.

References

[1] H. Zimmerman, "The OSI Reference Model—The ISO Model of Architecture for Open Systems Interconnection," *IEEE Transactions on Communications*, Volume 28, No. 4, April 1980, pp. 425–432.

—Craig Partridge, BBN Systems and Technologies

Write to *ConneXions*!

Have a question about your subscription? Are you moving, and need to give us your new address? Suggestions for topics? Want to write an article? A letter to the Editor? Have a question for an author? Need a *ConneXions* binder? Want to enquire about back issues? (there are now 74 to choose from; ask for our free 1987–1992 index booklet). We want to hear from you. Contact us at:

ConneXions—The Interoperability Report
480 San Antonio Road
Mountain View, CA 94040–1219
USA
Phone: +1 415-941-3399 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-949-1779
E-mail: connexions@interop.com

CONNEXIONS

Subscribe to CONNEXIONS

U.S./Canada ❑ $150. for 12 issues/year ❑ $270. for 24 issues/two years ❑ $360. for 36 issues/three years

International $ 50. additional per year (**Please apply to all of the above.**)

Name _____ Title _____
Company _____
Address _____
City _____ State _____ Zip _____
Country _____ Telephone ( ) _____

❑ Check enclosed (in U.S. dollars made payable to **CONNEXIONS**).
❑ Visa ❑ MasterCard ❑ American Express ❑ Diners Club Card # _____ Exp.Date _____
Signature _____

Please return this application with payment to: **CONNEXIONS**

Back issues available upon request $15./each
Volume discounts available upon request

480 San Antonio Road, Suite 100
Mountain View, CA 94040 U.S.A.
415-941-3399 FAX: 415-949-1779
connexions@interop.com